

The Cryptographic Implications of the LinkedIn Data Breach

Aditya Gune
gunea@oregonstate.edu

Department of Electrical Engineering and Computer Science
Oregon State University
Corvallis, OR 97333

March 19, 2017

Abstract

Data security and personal privacy are difficult to maintain in the Internet age. In 2012, professional networking site LinkedIn suffered a breach, compromising the login of over 100 million accounts. The passwords were cracked and sold online, exposing the authentication credentials millions of users. This manuscript dissects the cryptographic failures implicated in the breach, and explores more secure methods of storing passwords.

1 Introduction

LinkedIn is one of the largest social media and professional networking sites in the world. With over 400 million members as of 2017[10], any data leak risks the personal and professional information of millions of people. In 2012, LinkedIn's servers were breached, exposing the hashed passwords of approximately 117 million accounts over a course of several years.[15] These passwords were quickly decoded and sold online, indicating that LinkedIn's method of storing passwords was cryptographically insecure.

1.1 The password problem

The problem of storing passwords securely has existed for decades. Storing passwords in plaintext is clearly insecure, posing an immediate and uncontrollable risk in the event of a breach. Storing encrypted passwords also poses a risk due to its two-way nature – anything encrypted can also be decrypted; in addition, encryption keys are often stored on the same servers as the data being encrypted, meaning a breach is potentially disastrous.

Most websites opt for storing hashed passwords. The intent of storing hashed passwords is to provide a fail-safe to prevent identity theft in the event of a server breach. The method is intended to provide authentication *and* secrecy, and to prevent adversaries from being able to guess the passwords being hashed.

1.2 LinkedIn breach

Like most web services, LinkedIn hashed its passwords. The company passed user information through a SHA-1 hash function.[6] The original breach, speculated to be through a SQL injection attack, occurred on June 5th 2012, and was reported by a number of news agencies. LinkedIn confirmed the breach the following day. [17] It was first thought that the breach exposed around 6.5 million passwords. However, in 2016, the full dump of the hack was posted, exposing the accounts of over 117 million users, whose information had been compromised but not made public. [5] The severity of the breach was compounded by improper use of cryptographic hash functions to conceal the password plaintext.

2 Cryptography of Password Hashing

Password hashing works by taking a password input of variable length, and providing a fixed-length output that seems like a random string. SHA-1, the specific hash function used by LinkedIn, was published in the 1990s. SHA-1 can encode an up to 2^{64} bit input into a 160 bit message digest.[4] However, LinkedIn's implementation of SHA-1 failed to include salts – random numbers unique to every user – instead simply storing hashed passwords directly to the server.[8]

SHA-1 works by repeatedly padding the message and breaking it into 512-bit blocks, then iteratively passing these through a series of logical functions. Further information can be found in the official Secure Hash Standard publication FIPS 180-4.[12]

3 Analysis of Failure

Despite the illusion of secrecy provided by SHA-1, LinkedIn's method of storing passwords was dangerously insecure. Using an unsalted hash leaves the system vulnerable to rainbow table attacks.

3.1 Rainbow tables

Rainbow table attacks assemble a pre-compiled list of passwords into data structures known as *hash chains*. The hashes of these passwords have already been calculated[1]. However, the storage cost for a simple associative array of passwords to hashes would be very high. Rainbow tables use *reduction functions*,

which map individual hashes back to plaintext passwords. This creates a time-memory tradeoff, sacrificing time at the cost of memory.[13] By applying alternating sequences of hashes and reduction functions, an attacker can create chains of hashes and plaintexts:[13]

$$P_i \xrightarrow{\text{Hash}(P)} C_i \xrightarrow{\text{Red}(C)} P_{i+1} \cdots P_n$$

Figure 1: A schematic of a hash chain

Only the first and last elements in the chain – P_0 and P_n – are stored in memory, thereby greatly reducing the amount of memory required to store a large number of password - hash pairs. Without a salt, an adversary could apply the combination of reduction and hashing functions until they reached an endpoint on a hash chain, then take the corresponding starting plaintext and follow its hash chain until the target password was found.[13] This vulnerability is compounded by the fact that poor password choices are often repeated by multiple users. Security firm KoreLogic’s dump of the most common passwords exposed in the LinkedIn attack revealed that over 1 million users used the phrase “123456” as a password.[2]

The introduction of a salt drastically increases the computational cost required for an effective rainbow table. For example, the password ‘123456’ always hashes to the same value, meaning there is only one hash chain entry that needs to be stored for this password. However, if we add a 2-bit salt to the beginning of the password, we see that the adversary now has 2^2 more possible values, extending the length of the hash chains within the rainbow table:

```
$ hashlib.sha1(m)
('5a44cf4f2b0f2bfc7da6f386481f6afbc8aff73f', '01123456')
```

```
$ hashlib.sha1(s + m)
('5a44cf4f2b0f2bfc7da6f386481f6afbc8aff73f', '01123456')
('ac0e191df76d3714cb4e2c2659d51753775662d6', '10123456')
('3cf621ead5cc3885a4a5caef840aad7404bdee81', '11123456')
('b388959b842429b18180899f7b101cf7ed8667db', '00123456')
```

A sufficiently large salt (such as 128 bits) would require the adversary to store or compute 2^{128} more values for *each* password, making the cost of such a rainbow table prohibitively high.

4 Potential Solutions

Of the most common ways to store passwords, LinkedIn chose perhaps the least secure (aside from storing just plaintexts). Nonetheless, storing passwords in

a cryptographically secure way remains a difficult task for system administrators, who often resort to using non-cryptographically secure constructions or hash functions. Principally, cryptographic hash functions should account for the *cost* of computing a hashed password, and ensure that cost scales with faster microprocessors, a requirement that traditional hash functions often do not satisfy. UNIX’s *crypt* function was first introduced in 1976, and could hash approximately four passwords per second on contemporary hardware. By 1999, this had increased to 200,000 *crypt* operations per second. [16] Using GPU-driven parallel computing, password crackers were able to compute over 63 billion SHA-1 hashes per second in 2012. [7] This trend underscores a fundamental requirement of cryptographic hash functions: they must be resistant to brute-force attacks.

To accomplish this, a password hashing function must ensure that the computational cost of hashing a password is large enough that an adversary attempting to crack thousands or millions of passwords through brute force will find the exercise too time-consuming.

4.1 Key Derivation Functions

A key derivation function (KDF) takes a non-random source of input - such as a password - and derives a cryptographically strong secret key from it. Many contemporary systems for securing passwords use KDFs rather than standard hash functions. A secure KDF will output a key that is indistinguishable from a random string of the same length, and will draw from a uniform distribution of such strings. Most KDFs are constructed from two modules: one that “extracts” a pseudorandom key K from the source input, and one that uses K as the seed for a pseudorandom function (PRF) to produce several cryptographically secure pseudorandom keys.[9]

However, a KDF in itself is not a cryptographically secure way to store passwords, and may not be resistant to brute-force attacks. A KDF construction used for password hashing should increase the cost required to compute a password hash. The implementations discussed here – Bcrypt and scrypt – accomplish this through a method known as *key stretching*. Key stretching involves using a KDF that requires 2^k cryptographic operations per hash computation, regardless of the length of the password. For example, an n -bit password would require 2^{k+n} operations under this system. [14]

4.2 Bcrypt

Bcrypt is an adaptive key derivation function, and is perhaps one of the most widely recommended methods for storing passwords securely. As discussed above, password hashing functions must be resistant to brute force attacks. Bcrypt provides this resistance by allowing administrators to specify the number of internal iterations required for a password hash calculation. This causes the computational cost of hashing a password, making brute force attacks prohibitively expensive. Bcrypt is based on the Blowfish cipher, a “64-bit block

cipher structured as a 16-round Feistel network.” [16]

$\begin{aligned} &\text{BLOWFISH}(m):[16] \\ &L_0 := m[: half] \\ &R_0 := m[half :] \\ &\text{for } i = 1 \text{ to } 16: \\ &\quad R_i := L_{i-1} \oplus P_i \\ &\quad L_i := R_{i-1} \oplus F(R_i) \\ &R_{17} := L_{16} \oplus P_{16} \\ &L_{17} := R_{16} \oplus P_{18} \\ &\text{return } L_{17} R_{17} \end{aligned}$

Though reduced round Blowfish does have vulnerabilities, the standard Blowfish block cipher has been demonstrated secure. [3] The Bcrypt algorithm uses the variant *Eksblowfish*, which takes in the user’s desired cost of computation as a parameter:

$\begin{aligned} &\text{EKSBLOWFISH}(cost, salt, key):[16] \\ &state \leftarrow \text{InitState}() \\ &state \leftarrow \text{ExpandKey}(state, salt, key) \\ &\text{repeat } 2^{cost}: \\ &\quad state \leftarrow \text{ExpandKey}(state, 0, salt) \\ &\quad state \leftarrow \text{ExpandKey}(state, 0, key) \\ &\text{return } state \end{aligned}$
--

Bcrypt is designed to remain resistant to brute force attacks regardless of the speed of modern microprocessors. The ability of Bcrypt to iterate for 2^{cost} cycles allows administrators to increase the number of iterations to compensate for stronger password cracking hardware. As the number of iterations increases, individual password hashes take longer to compute, ensuring that a brute force attack remains an expensive task for adversaries. [11]

4.3 Scrypt

Scrypt follows a similar paradigm to Bcrypt, using an adaptive cost mechanism to stymie brute force attackers. However, Bcrypt is still vulnerable to brute force attacks using parallelized hardware. Improvements in semiconductor technology allow attackers to embed more circuits at lower cost. [14]

Highly parallelized hardware such as the GPU is well suited to the repetitive task of password cracking, because the large number of circuits allows even computationally expensive passwords to be hashed non-sequentially. Billions of password hashing operations can be performed every second, greatly reducing the effectiveness of even adaptive KDFs like Bcrypt. [7]

Scrypt provides resistance against parallelized brute force attacks by making both sides of the time-memory tradeoff costly. It allows the amount of memory required to hash a password to increase proportionally to the number of computations.

$\text{MFCRYPT}(\textit{password}, \textit{salt}, N, p, dkLen):[14]$ $(B_0 \cdots B_{p-1}) \leftarrow \text{PBKDF2}(P, S, 1, p * MFLen)$ for $i = 0$ to $p - 1$: $B_i \leftarrow MF(B_i, N)$ $DK \leftarrow \text{PBKDF2}_{PRF}(P, B_0 B_1 \cdots B_{p-1}, 1, dkLen)$

The full proof of this algorithm can be found in [14]. Scrypt’s large memory requirement is partly due to the action of computing and storing in memory a large number of bit strings using a pseudorandom function. This requires allocating a large number of memory locations. The amount of memory used and number of computations required can both be tuned using the parameters N and p . This allows system administrators to ensure that the cost of hashing millions of passwords remains very expensive (both in time and in memory), even with advancements in hardware.

5 Conclusion

The use of cryptographically secure password storage remains critical to information security as a whole. Password hashing is a last line of defense, protecting user data in the event of a breach. The 2012 LinkedIn breach is an excellent case study in securing stored passwords.

The use of an unsalted hash function led to the exposure of over 100 million passwords. As we have shown, standard hash functions such as SHA-1 are no longer adequate for storing passwords, and are vulnerable to both rainbow table and brute force attacks. Alternatives such as Bcrypt and scrypt use combinations of salts and adaptive computation and memory cost, to make password cracking more expensive than what a standard unsalted hash function would allow.

References

- [1] CAPEC-55: Rainbow Table Password Cracking, 2015.
- [2] LinkedIn Revisited - Full 2012 Hash Dump Analysis, 2016.
- [3] ALABAICHI, A., AHMAD, F., AND MAHMOD, R. Security analysis of blowfish algorithm. *Informatics and Applications, Second International Conference* (2013).
- [4] EASTLAKE, D., AND JONES, P. US Secure Hash Algorithm 1 (SHA1). RFC 3174, RFC Editor, September 2001.
- [5] FRANCESCHI-BICCHERAI, L. Another Day, Another Hack: 117 Million LinkedIn Emails And Passwords, 2016.
- [6] GOLDMAN, D. More than 6 million LinkedIn passwords stolen, 2012.

- [7] GOSNEY, J. Password Cracking HPC, 2012.
- [8] GOSNEY, J. How LinkedIn’s password sloppiness hurts us all, 2016.
- [9] KRAWCZYK, H. Cryptographic Extraction and Key Derivation: The HKDF Scheme. *International Association of Cryptologic Research* (2010).
- [10] LINKEDIN. About us, 2017.
- [11] MALVONI, K., DESIGNER, S., AND KNEZOVIC, J. Are Your Passwords Safe: Energy-Efficient Bcrypt Cracking with Low-Cost Parallel Hardware. *USENIX* (2014).
- [12] NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY. FIPS PUB 180-4: Secure Hash Standard, 2014.
- [13] OECHSLIN, P. Making a Faster Cryptanalytic Time-Memory Trade-Off. *Annual International Cryptology Conference* (2003).
- [14] PERCIVAL, C. Stronger Key Derivation via Sequential Memory-Hard Functions. *TarSnap* (2009).
- [15] PEREZ, S. 117 million LinkedIn emails and passwords from a 2012 hack just got posted online, 2016.
- [16] PROVOS, N., AND MAZIÈRES, D. A Future-Adaptable Password Scheme. *USENIX* (1999).
- [17] SILVEIRA, V. An Update on LinkedIn Member Passwords Compromised, 2012.